arXiv:1404.2910v2 [math.ST] 11 May 2015

STATISTICAL TESTS FOR LARGE TREE-STRUCTURED DATA

KARTHIK BHARATH, PRABHANJAN KAMBADUR, DIPAK. K. DEY, ARVIND RAO,
AND VEERABHADRAN BALADANDAYUTHAPANI

Abstract. We develop a general statistical framework for analyzing large tree-structured data with a focus on developing asymptotic goodness-of-fit tests. The tests use the distributional properties of the Continuum Random Tree (CRT), introduced by Aldous [1993], which arises as the invariant limit for a broad class of models for tree-structured data based on conditioned Galton–Watson processes. On the basis of two characterizations of the CRT that involve a special class of subtrees and Brownian excursion, we evaluate its utility in distinguishing between populations of trees. The proposed generative probability models for tree-structured data are easily interpretable and lead to an efficient scheme for simulating a large number of trees with arbitrary sizes. The resulting test statistics for the goodness-of-fit tests are simple to compute and are asymptotically distributed as χ^2 and F random variables, allowing for a thorough examination of their properties. We illustrate our methods on an important application of detecting tumour heterogeneity in brain cancers. A novel approach with tree-based representations of magnetic resonance images is used wherein the developed tests are employed to ascertain tumour heterogeneity between two groups of patients.

Keywords: Brownian excursion; Conditioned Galton–Watson trees; Dyck path; Goodness-of-fit tests.

1. Introduction

Rapid advancements in technologies has led to the emergence of datasets wherein the underlying quantities of interest are non-Euclidean objects. A particular example, increasingly encountered across several disciplines, is data that generate tree-structured data or allow for tree-like, hierarchical representations. These include, for example, database applications that involve the detection of duplicate XML records [Yang et al., 2005, Tatikonda and Parthasarathy, 2010]; modelling secondary structures of RNA sequences [Schmitt and Waterman, 1994, Le et al., 1989, Kolbe and Eddy, 2009]; modelling of observed blood vessels in a human brain [Aydin et al., 2011, Shen et al., 2013, Wang and Marron, 2007], and tasks of protein classification [Busch et al., 2009]. Following a suitable representation of trees, the assumption then is to treat the tree-structured object as an observed or an observable quantity that represents the statistical atom.

Tree-structured data can be regarded as data observed on acyclic graphs with a distinguished vertex that is referred to as the root, and as a consequence, can be represented in several ways, with the choice of representation usually dictated by the application and ease of computation. These include an adjacency matrix representation that captures only topological information, set-theoretic representation of a tree through its vertex and branch-length sets, and coding a tree through a bijective step function, which is popular in combinatorics. However, some central challenges have stymied the systematic development of tools for statistical inference: the non-Euclidean nature of the underlying space introduces an unreasonable dependence of probability models on the choice of

representation and geometry; tree-structured data rarely contain the same number of vertices, which leads to issues when comparing trees of differing sizes; and it is not straightforward to sample trees from a purported probability model that accounts for topological and branch length information.

Attempts to address these issues, at least from a modelling perspective, have hitherto been characterised by nonparametric or heuristic/algorithmic approaches specifically motivated by applications. Shen et al. [2013] used nonparametric functional data methods on the bijective representation of a tree as a continuous function, known as Dyck paths, using an aligning mechanism for exploring modes of variation in the trees in regression settings; however, this leads to the creation of some spurious artifacts for tree-structures, for example, negative branch lengths. Aydin et al. [2009], Aylward and Bullitt [2002] modelled the brain vessel data as three-dimensional trees embedded on a plane with the objective of investigating the relationship between the brain vessel structure and an individual's age. The Edit distance coupled with algorithmic tools was used by Yang et al. [2005] and Tatikonda and Parthasarathy [2010] to detect duplicate XML records; they represented each document that contained a hierarchy of program statements and tags as a rooted tree and used a collection of such documents to generate a sample of tree-structured data. Another important example of the use of algorithmic methods is to represent the observed secondary structures of RNA sequences by means of ordered trees, wherein changes in RNA sequences are described as suitable transformations of trees [Schmitt and Waterman, 1994, Le et al., 1989, Kolbe and Eddy, 2009]. Despite the flexibility associated with such nonparametric and algorithmic methods for tree-structured data, depending on the choice of representation and metric, developing robust inferential methods is not straightforward, and determining distributions of the statistics of interests is difficult. Critically, without a probability model supporting the simulation of trees through a generative model, it becomes hard to assess the generality of the methods across tree-structured datasets.

In contrast, the development of parametric probability models for trees was considered by Steele [1987], Aldous [1996] wherein the parameters were designed to capture topological features and branch length information, making possible parametric inference on tree-structured data. Such probability models form the basis of modelling hierarchical data using Bayesian methods with nonparametric hierarchical priors such as the Dirichlet diffusion tree [Neal, 2001], beta-coalescent trees [Hu et al., 2013] and stick-breaking priors [Adams et al., 2010], and tasks involving clustering [Heller and Ghahramani, 2005, Teh et al., 2008]. These prior distributions are basically distributions over multivariate distributions that capture hierarchical relationships. Although it is possible to use such methods to simulate the required classes of trees, inference depends on the ability to accurately generate samples of trees from a generally intractable posterior distribution on trees; assessment of the quality of the posterior samples is particularly hard since the sample space is non-standard, and is generally not well understood. Furthermore, the nature of inferential methods dictates that extensions to general classes of trees require substantial revamping of the computational setup; for example, modifying the Dirichlet diffusion tree prior for binary trees to handle trees with an arbitrary number of children at each node requires significant changes in the probability model and computational techniques.

Such considerations point towards an amalgamation of the simplicity, in terms of interpretation and simulation, of parametric generative models, and the flexibility and generality offered by non-parametric or algorithmic methods, towards developing coherent and principled inferential methods for tree-structured data. This forms the *leitmotif* of this article. What is crucial, considering the complexity of the data, is that one is able to systematically assess the advantages and shortcomings of the proposed model in view of the application of interest. In order to tackle such a requirement, we seek recourse to the asymptotic theory of random trees through the the abstract notion of a *Continuum Random Tree*(CRT) from Aldous [1991a] and Aldous [1993]. The CRT arises as the weak invariant continuous limit as vertices grow without bound for a large class of random trees modelled by critical Galton–Watson processes conditioned on total progeny, known as *conditioned Galton–Watson* processes.

The main aim of this article is two-fold:

(1) to provide simple but general generative models for tree-structured data based on conditioned Galton–Watson trees and the CRT, along with an efficient simulation mechanism;
(2) to demonstrate the utility of such models in developing goodness-of-fit tests, based on χ^2 (one-sample) and F (two-sample) test statistics, for tree-structured data.

Starting with a simple model for binary trees that leads to a testing procedure, we show how the model can be extended to a parametric model for general trees through the CRT using conditioned Galton–Watson tree models, and construct goodness-of-fit tests. The probability models via conditioned Galton–Watson processes are applicable to a large class of trees, and can be defined with a canonical parameter, which in this case will be the variance of the offspring distribution. As a consequence, we are able to design a simulation technique based on an expected linear run-time algorithm by Devroye [2012], which enables efficient generation of large samples of large trees. The price we pay for the above-mentioned attractive properties is that the framework is asymptotic in spirit and useful primarily when the trees contain a large number of vertices. The weak convergence to CRT of conditioned Galton–Watson processes is an invariance principle of sorts, i.e., *every* critical conditioned Galton–Watson process (up to a proportionality constant), regardless of the offspring distribution, converges to the CRT. Exploiting this property along with two different distributional characterisations of the CRT, we propose goodness-of-fit tests for a broad class of tree-structured data based on familiar χ^2 and F test statistics.

We illustrate our methodology through a case study in cancer imaging. It is well-established that a systematic classification of intra- and inter-tumour heterogeneity is crucial for drug development and for accurate assessment of response to treatment, especially in cancer. We consider the problem of detecting brain tumour heterogeneity between two groups of patients from magnetic resonance (MR) images. Tumour heterogeneity expressed though voxel intensities indicates a latent ordering of groups of voxels with similar intensities that represent similar etiologies. Current approaches to this problem are based on simple summaries of the entire image such as skewness or kurtosis of the distribution of the intensities; see Just [2014] for a detailed review. While potentially revealing, these approaches fail to take into account structural complexities of the pixel-level intensities. We take an alternative modelling strategy, wherein we use tree-based representations of the image obtained with hierarchical clustering. The properties of the tree, such as the branching structure and distances

between groups of pixels (of functions of it), serve as natural summary statistics for evaluating tumour heterogeneity. Our approach is to represent the hierarchical structure as a dendrogram or a strict binary tree, consider a statistic that relates the pixel-cluster distances to the branch lengths of a carefully chosen subtree of the dendrogram, and construct tests for heterogeneity based on this statistic. Asymptotically, the conditioned Galton–Watson model, through its connections with the CRT, helps us determine the distribution of this test statistic and conduct rigorous statistical inference.

The paper is structured as follows. In Section 2, we look at a simple model for binary trees that leads to the definition of a density for the tree-structured data, and the development of tests for tree distributions. The simplicity of this model motivates the use of conditioned Galton–Watson tree models through the CRT. In Section 3, we review the key ingredients of the CRT, including conditioned Galton–Watson trees, Dyck paths and least common ancestors trees, restricting our attention to the aspects of their relationships that are relevant to our purposes. In Section 4, we consider the issue of estimating and interpreting the variance parameter of the offspring distribution of conditioned Galton–Watson tree models. Then in Section 5, we construct goodness-of-fit tests for conditioned Galton–Watson tree models based on two related approaches. In Section 6, we conduct simulations by generating a large number of very large conditioned Galton–Watson trees using the expected linear-time algorithm provided by Devroye [2012], and examine the performances of the proposed tests. In Section 7, we implement the tests on a dataset of brain tumour images using a tree representation to check for group differences between cancer patients who experience long versus short survival times. In Section 8, we comment on possible extensions and generalizations, and point out some shortcomings of the current framework. Proofs of the results are relegated to Section 9. Computation and simulation details are available in the Supplementary Material.

2. A generative model and test for binary trees

Throughout this article, we confine our attention to finite, rooted, ordered trees: trees with a distinct vertex, referred to as the root, containing a finite number of vertices, and having an ordering amongst the offspring at a parent node. This leads to the notion of a left and a right child in the case of binary trees. However, we allow the trees to have non-unit branch lengths and unequal numbers of vertices.

We begin by describing a simple generative model for strict binary trees, where every internal vertex has 0 or 2 children, and the resulting goodness-of-fit tests. To avoid confusion, throughout the article, when we refer to binary trees, we mean strict binary trees. The model provides a motivating reason for using the CRT for extension to trees that are not necessarily binary, and is used in the construction of the CRT [Aldous, 1993]. The main benefit of this construction is that a coherent probability density can be assigned to a binary tree. Suppose $t_1, t_2, \ldots$ are the arrival times of a non-homogeneous Poisson process with rate $\lambda(t) = t$. Let $\tau(1)$ be a tree with one leaf formed by attaching an edge of length t_1 to the root. This tree will have a density that corresponds to the time till the first event of the non-homogeneous Poisson process, with mean function $m(t) = \int_0^t s ds = t^2/2$. A tree with two leaves $\tau(2)$ is then constructed by adding a second edge of length $t_2 - t_1$, the inter-arrival time, to a point on $\tau(1)$ chosen according to a uniform

distribution on $[0, t_1]$. Recursively, the tree $\tau(k+1)$ is obtained from $\tau(k)$ by attaching an edge of length $t_{k+1} - t_k$ to a point chosen randomly on the tree—i.e., according to a uniform distribution on the sum of the branch lengths of $\tau(k)$–and labelling the new leaf $k+1$. Since we are interested in only ordered trees, an edge can be attached to the left or to the right of the randomly chosen point with equal probability. Observe that $\tau(k)$ has $2k-1$ branches with lengths, say $(x_1, \ldots, x_{2k-1})$, and $\tau(k+1)$ has $2k+1$ branches of lengths $(y_1, \ldots, y_{2k+1})$, with $\tau(k+1)$ formed by the above construction by splitting a branch of length x_j into two branches y_{j1} and y_{j2} with $x_j = y_{j1} + y_{j2}$. Therefore, through induction, the density of $\tau(k+1)$ is

$$f(\tau(k+1)) = f(\tau(k)) \times (\text{density of } k+1\text{th inter-arrival time})$$
$$\times (\text{probability of the attaching new edge to } \tau(k)) \times (\text{probability of left or right attachment}).$$

Denoting the sum of the branch lengths of $\tau(k)$ by s_k and noting from the construction that the third term is $1/s_k$ and the fourth is $1/2$, we obtain

$$f(\tau(k+1)) = f(\tau(k)) s_{k+1} e^{-\frac{1}{2}(s_{k+1}^2 - s_k^2)} \frac{1}{2s_k},$$

which reduces to

$$f(\tau(k)) = \frac{1}{2^{k-1}} s e^{-s^2/2}, \quad s = \sum_{i=1}^{2k-1} x_i.$$

The function f can be interpreted as assigning mass to a tree with branch lengths $(x_1, \ldots, x_{2k-1})$ with a canonical binary tree topology on k leaves or $2k$ vertices. That is, the mass is the same for rooted binary trees with k leaves of all possible topologies with branch lengths $(x_1, \ldots, x_{2k-1})$. The function $f(\cdot)$ is non-negative but does not integrate to one. Noting, from Aldous [1993], that

$$\int_{x_1} \cdots \int_{x_{2k-1}} s e^{-s^2/2} dx_1 \ldots dx_{2k-1} = \prod_{i=1}^{k-1} \frac{1}{2i-1},$$

the probability density, with respect to the Lebesgue measure on $\mathbb{R}_+^{2k-1}$, of the ordered and rooted binary tree with k leaves is

$$f(\tau(k)) = \left[\prod_{i=1}^{k-1} \frac{1}{2i-1} \right]^{-1} \frac{1}{2^{k-1}} s e^{-s^2/2}, \quad s = \sum_{i=1}^{2k-1} x_i. \tag{1}$$

Two properties of interest arise when examining the density f: first, f is independent of the topology of the tree, and permutation of the leaves does not alter the density, which implies exchangeability of the branch lengths; and second, the density f *depends on the branch lengths only through their sum*. Exchangeability of leaves is particularly attractive when considering binary trees obtained from hierarchical clustering wherein the data reside at the leaves; this will be examined in detail in the data application. In a certain sense, the topology of the tree is captured in the product term acting as the normalising factor and in the fact that s is made up from k leaves, from which any of the $\prod_{i=1}^{k-1}(2i-1)$ possible binary trees can be constructed. The second property implies that the density of the entire tree is captured through its total path length, which is the sum of the branch lengths $s = \sum_{i=1}^{2k-1} x_i$, or in terms of the process, the kth arrival time.

From the simple model, a parametric form of the density can be obtained by defining $\lambda(t, \theta) = \theta t$, which leads to a parametric family $\{f_\theta; \theta \in \Theta\}$, wherein trees of varying branch lengths can be generated by varying θ in Θ. This leads to the definition of a generative model which can be used to equip the sample with a parametric likelihood. Simulation of such trees reduces to efficient simulation of a non-homogeneous Poisson process with a simple intensity function.

2.1. **One- and two-sample tests.** The total path length of the binary tree constructed from the non-homogeneous Poisson model can be used to construct distributional tests for the model that generates the binary trees, since they completely determine the density f once the number of leaves is known. It is pertinent to note that there is a difference between the total path length that characterizes f and the kth arrival time that characterizes the process. Recall that while constructing $\tau(k+1)$ from $\tau(k)$, the new edge with length given by the inter-arrival time is attached uniformly on $\tau(k)$ with total path length s; this action is key to obtaining the density in (1), and depends on the non-homogeneous Poisson process only through the time of the kth arrival. As a consequence, once the number of leaves is observed, the density of the tree arising from the construction is fully characterised by the total path length. Generally speaking, in contrast to the homogeneous Poisson process with a constant rate function, the distribution of the kth arrival time, in general, does not characterise a non-homogeneous Poisson process. However, it is always possible to convert a non-homogeneous Poisson process to a homogeneous one with unit rate by warping the time scale: in our setup, if $N(t)$ is the non-homogeneous Poisson process with intensity function $m(t) = t^2/2$ (rate $\lambda(t) = t$), then the process $N(m^{-1}(t)) = N(\sqrt{2t})$ is a homogeneous Poisson process with a rate of one. Under the non-linear time scale of the homogeneous Poisson process, in order to obtain a density that depends solely on the total path length, we need to to appropriately modify the step in the construction that involves choosing a point on $\tau(k)$ at which to attach a new edge to obtain $\tau(k+1)$. The following proposition determines the distribution of the total path length.

Proposition 1. Suppose $\tau(k)$ is a binary tree with k leaves and branch lengths $x_1, \ldots, x_{2k-1}$ generated from the non-homogeneous Poisson model. Then the total path length $\sum_{i=1}^{2k-1} x_i$ follows a Gamma distribution with shape k and scale 2.

Using Proposition 1, we obtain exact one- and two-sample tests for binary trees from the non-homogeneous Poisson process model of any size, which is summarised in the following result and for which the proof is straightforward.

Theorem 1. *Suppose* $\boldsymbol{\tau(n)} = (\tau(n_1), \ldots, \tau(n_p))$ *and* $\boldsymbol{\eta(m)} = (\eta(m_1), \ldots, \eta(m_q))$ *are independent samples of binary trees from* π_τ *and* π_η *respectively.*

(1) Consider the critical function

$$\phi(\mathbf{n}, \alpha) = \begin{cases} 1 & \text{if } \sum_{i=1}^{p} s_i > \chi_{1-\alpha, 2\sum_{i=1}^{p} n_i}; \\ 0 & \text{otherwise,} \end{cases}$$

where s_i is the sum of the branch lengths of $\tau(n_i)$ and $\chi_{\alpha,b}$ denotes the αth percentile of a χ^2 distribution with b degrees of freedom. For the hypotheses $H_0 : \pi_\tau = f$ against $H_1 : \pi_\tau \neq f$, where f is the density in (1), $E_{H_0}\phi(\mathbf{n}, \alpha) = \alpha$.

(2) Let r_j denote the sum of the branch lengths of $\eta(m_j)$, and without loss of generality assume that $\sum_{i=1}^{p} s_i > \sum_{j=1}^{q} r_j$. Then, the critical function

$$\psi(\mathbf{n}, \mathbf{m}, \alpha) = \begin{cases} 1 & \text{if } \frac{\sum_{i=1}^{p} s_i}{\sum_{j=1}^{q} r_j} > \left(\frac{\sum_{i=1}^{p} n_i}{\sum_{j=1}^{q} m_j}\right) F_{1-\alpha, 2\sum_{i=1}^{p} n_i, 2\sum_{j=1}^{q} m_j}; \\ 0 & \text{otherwise,} \end{cases}$$

where $F_{\alpha,a,b}$ is the αth percentile of an F distribution with a and b degrees of freedom, for testing $H_0 : \pi_\tau = \pi_\eta = f$, is such that $E_{H_0}\psi(\mathbf{n}, \mathbf{m}, \alpha) = \alpha$.

Except for the number of leaves in each tree and hence in the sample, the distribution of the test statistic used for the tests is independent of the topology of the binary tree. Nonetheless, the simplicity of the tests is particularly attractive, and the performance of the tests can be assessed in detail. Applications wherein the branch length information in binary trees is important would provide an appropriate setting for the use of these tests.

3. Preliminaries: Conditioned Galton–Watson trees and the CRT

The model from Section 2, albeit simple and tractable, is not flexible enough to capture topological information about trees. Moreover, the tests are restricted to the generative model, and cannot be viewed as a goodness-of-fit test for trees. However, without eschewing the simplicity, we wish to extend the use of the model to more general trees, and develop goodness-of-fit tests. In this section we describe in considerable detail the ingredients required to extend the use of the density in (1) to non-binary trees as well. Starting with conditioned Galton–Watson process models for trees, we briefly describe the convergence to CRT, and detail two of its characterisations, one of which leads to the density in (1). It will be seen that the density in (1) is the probability density of a certain class of subtrees of the CRT that completely characterises the law of the CRT.

Following the notation used by Aldous [1991a], we represent a finite rooted tree τ_n with n vertices, including the root, as set of vertices $\mathcal{V}(\tau_n) = (\text{root}, v_1, \ldots, v_{n-1})$ and a set of edges $\mathcal{E}(\tau_n) = (e_1, \ldots, e_{n-1})$, represented as a point

$$\tau_n = \Big(\mathcal{V}(\tau_n), \mathcal{E}(\tau_n)\Big)$$

in the space $\mathcal{T}_n \times \mathbb{R}_+^{n-1}$ where $\mathcal{T}_n$ is the set of all finite trees on n vertices. Note the difference in notation to $\tau(n)$ which is the tree with n terminal vertices or leaves.

3.1. Conditioned Galton–Watson trees.

A Galton–Watson tree is defined in the following manner: given a probability distribution $(\pi_k, k = 0, 1, \ldots)$ on the non-negative integers, or equivalently a random variable ξ with distribution π_k, we construct a Galton–Watson tree τ recursively starting with the root and giving each node a number of children that is an independent copy of ξ; $P(\xi = k) = \pi_k$ for $k = 0, 1, \ldots$ is referred to as the offspring distribution and the out-degrees of the

vertices are i.i.d. copies of ξ from π_k. As a consequence, ξ induces a unique distribution on τ as

$$P(\tau = t) = \prod_{v \in \mathcal{V}(t)} \pi_{o(v,t)},$$

where $o(v,t)$ is the out-degree or the number of children of vertex v in the vertex set $\mathcal{V}(t)$ of tree t. If one wishes to model a set of tree-structured data using Galton–Watson trees, two critical issues arise at this point: (a) Galton–Watson trees, with positive probability, can be infinite, whereas observed trees in practice are always finite; and (b) the observed trees must have been randomly sampled from an appropriate population. How do we ensure random sampling in an appropriate population? We address these issues by considering conditioned Galton–Watson trees, which are Galton–Watson trees conditioned on total progeny, known as simply generated trees in combinatorial literature. That is, the distribution of a conditioned Galton–Watson tree τ_n conditioned to have n vertices is

$$P(\tau_n = t) \propto \prod_{v \in \mathcal{V}(t)} \pi_{o(v,t)} \quad \text{on } \{t : \text{ cardinality of } \mathcal{V}(t) = n\}.$$

Importantly, it is known that for a fixed offspring distribution π_k, the corresponding conditioned Galton–Watson tree can be viewed as being picked according to a uniform distribution on certain types of trees with n vertices. For example, if we wish to choose a binary tree with n vertices according to a uniform distribution on the space of n-vertex binary trees, then, we can equivalently construct a conditioned Galton–Watson tree with an offspring distribution of $1-p$ and p each for 0 and 2 children, respectively, where p is the probability of having 2 children. Conditioned Galton–Watson trees comprise a large class of trees that offer flexibility from a modelling perspective. We enumerate a few useful classes that offer a broad range of models for tree-structured data based on specific parameterizations.

(i) Ordered trees with unrestricted degree: conditioned Galton–Watson trees with the offspring distribution given by Geo(1/2), a geometric distribution with success probability 1/2.

(ii) Binary trees (not strict): conditioned Galton–Watson trees with the offspring distribution Bin(2,1/2), where vertices contain 0,1 or 2 children according to a binomial distribution, with 2 trials and success probability 1/2.

(iii) Strict binary trees that are ordered: conditioned Galton–Watson trees where the vertices contain either 0 or 2 children with equal probability.

(iv) Unary-binary trees that are ordered: conditioned Galton–Watson trees with the offspring distribution UB(1/3,1/3), where the vertices contain 0, 1 or 2 children with equal probability.

(v) Unary-binary trees that are unordered and unlabelled: conditioned Galton–Watson trees with vertices containing 0, 1 or 2 with probabilities $\pi_0 = \frac{1}{2+\sqrt{2}}$, $\pi_1 = \frac{\sqrt{2}}{2+\sqrt{2}}$ and $\pi_3 = \frac{1}{2+\sqrt{2}}$, respectively.

(vi) m-ary trees: conditioned Galton–Watson trees with vertices that contain $0, 1, \ldots, m$ for $m \geq 3$ children with the distribution given by a binomial with m trials and success probability $1/m$.

Note that the enumerated offspring distributions can be parameterized by the variance parameter. This will turn out to be useful when developing simulation schemes. Plainly, there is no branch length information in the conditioned Galton–Watson tree models. The offspring distributions above are all of unit mean, implying that the trees were constructed from a critical Galton–Watson

process. The criticality of the process turns out to be a sufficient condition for the corresponding conditioned Galton–Watson trees to converge to the limit CRT; however, for modelling purposes we require more. For example, it is not enough that the only potential model for strict binary trees is the conditioned Galton–Watson tree with offspring distribution 1/2 for both 0 and 2 children; similar arguments apply to the other cases. The following proposition demonstrates the possibility of extending the class of potential offspring distributions to sub and super-critical cases for modelling.

Proposition 2. Conditioned Galton–Watson trees with the following offspring distributions can be modelled as critical conditioned Galton–Watson trees for every $0 < p < 1$ in (i)-(iii) and $0 < p_0, p_1 < 1$ in (iv).

(i) Geo(p): $\pi_i = (1-p)^{i-1}p$ for $i = 1, 2, \ldots$ and $0 < p < 1$;
(ii) Bin($2, p$): $\pi_i = \frac{2!}{(2-i)!i!}p^i(1-p)^{2-i}$ for $i = 0, 1, 2$ and $0 < p < 1$;
(iii) Bin(p): $\pi_0 = 1-p, \pi_2 = p$ and $0 < p < 1$;
(iv) UB(p_0, p_1): $\pi_0 = p_0, \pi_1 = p_1, \pi_2 = 1 - p_0 - p_1$ and $0 < p_0, p_1 < 1$.

Proposition clarifies the assumption of criticality for the convergence of conditioned Galton–Watson trees to the CRT: in the sub-critical or super-critical case with offspring distribution π_i , as long as there exists a $\lambda > 0$ such that $\sum_{i\geq 0}\pi_i\lambda^i < \infty$, the asymptotic behaviour resembles that of a critical Galton–Watson process with finite offspring variance. At this point we may very well choose to view inference on conditioned Galton–Watson trees as inference on Galton–Watson processes; a good source for available methods is Guttorp [1991]. However, we move to the asymptotic setting for two reasons: to incorporate branch length information within the conditioned Galton–Watson tree models; and to obtain knowledge of the distributions of local structural aspects like height and variations in branching structure, through weak convergence techniques based on the CRT. In this article, we shall consider conditioned Galton–Watson trees and refer to them simply as *trees* when the context is clear.

3.2. **Continuum Random Tree and its characterisations.** We give a heuristic definition of the CRT that suffices for our needs and refer the interested reader to Aldous [1993] for details. The CRT arises as the continuous limit of conditioned Galton–Watson trees when the number of vertices is allowed to grow without bound. One way to compare trees of different sizes is to embed $\tau_n = (\mathcal{V}(\tau_n), \mathcal{E}(\tau_n))$ as an element of the linear space ℓ_1, the Banach space of infinite sequences $x = (x_1, x_2, \ldots)$ such that $||x|| = \sum_i |x_i| < \infty$. Such an embedding makes it possible to compare and scale trees of different number of vertices in a natural way since the space affords an infinite number of "orthogonal" directions in which the branches can lie without the possibility of collisions with other parts of the tree, which is grow without bound. The CRT or the *Brownian* CRT is the unique tree that is the closure of the union $\cup_{n\geq 1}\tau_n$. Its distribution can be specified in several ways, of which the following two will be of primary concern in this article: as a weak limit of family trees of critical conditioned Galton–Watson trees; and distributions of a class of subtrees that are analogous to finite dimensional projections of a stochastic process. Both specifications are intimately connected to a standard Brownian excursion. We first need to review two constructs not restricted to conditioned Galton–Watson trees, which play a pivotal role in the characterisation of the CRT.

3.2.1. *Dyck paths.* Manipulating trees-structured data, theoretically and computationally, is difficult since topological and branch length information needs to stored and carried over in computations. As a consequence, a convenient alternative is to code a tree by a *Dyck path*, described as follows. Any rooted ordered tree on n vertices can be uniquely coded by a traversal of the tree; when the traversal is a depth-first walk, one can construct a function that is bijective to the tree in the following manner. For ease of exposition, assume that the edges or branches of a tree τ_n with n vertices have length 1. For a fixed positive integer n, Dyck paths are lattice excursions of length $2n$: sequences $(d_j, 0 \leq j \leq 2n)$, where $d_0 = d_{2n} = 0$ and $d_j > 0$ with $d_{j+1} - d_j \in \{-1, +1\}$ for all $0 \leq j \leq 2n-1$. Imagine the motion of a particle that starts at time $t = 0$ from the root of the tree and then explores the tree from the left to the right, moving continuously along the edges at *unit speed* until all the edges have been explored and the particle has come back to the root. Since it is clear that each edge will be crossed twice in this evolution, the total time needed to explore the tree is $2n$. The value $H_n(s)$ of a continuous function $H_n : [0, 2n] \to \mathbb{R}_{\geq 0}$ at time $s \in [0, 2n]$ is the distance (on the tree) between the position of the particle at time s and the root; figure 1, taken from Pitman [2006], offers a more intuitive description with edge lengths not all equal to one. When branch lengths are taken into consideration, the depth-first walk is performed at unit speed so that H_n always has slope 1 or -1. Therefore, if τ_n is a tree of size n, the sequence $(H_n(0), H_n(1), \cdots, H_n(2n))$ is its Dyck path of length $2n$. The representation of the Dyck path of a tree τ_n in terms of the distance d between its vertices is related to the function H_n as

$$H_n(s) = d(\text{root}, v), \tag{2}$$

where v is the vertex obtained during the depth-first walk such that the sum of the edges traversed till v is s, and $d(u, v)$ is the length of the unique path from vertex u to vertex v. This discussion is formalized by the following proposition, for which the proof is straightforward.

Proposition 3. The map $\tau_n \mapsto (H_n(0), H_n(1), \cdots, H_n(2n))$ is a bijection from the set of rooted ordered trees with size n to the set of all Dyck paths of length $2n$.

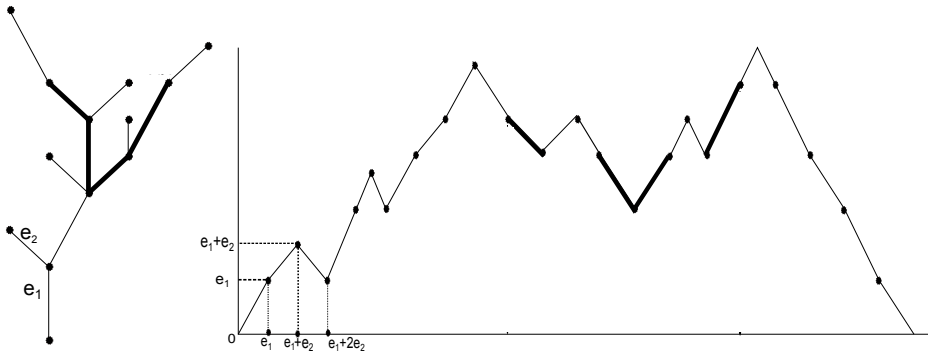

FIGURE 1. A tree with root at the bottom and its corresponding Dyck path. The x axis ranges from 0 to twice the sum of lengths of the edges; the Dyck path is constructed by traversing the tree in a depth-first manner at unit speed.

3.2.2. *Least common ancestor (LCA) subtrees.* Here, we define the class of spanning binary subtrees that characterise the distribution of the CRT. For a conditioned Galton–Watson finite tree $\tau_n = (\mathcal{V}(\tau_n), \mathcal{E}(\tau_n))$, we define its least common ancestor tree, or L-tree, in the following manner: choose a subset B of $\mathcal{V}(\tau_n)$; for vertices v_1 and v_2 in B find their last common ancestor, or the branch point after which the paths to the v_1 and v_2 from the root diverge or branch out. Now, the L-tree corresponding to the subset B of the vertices of τ_n is the tree, denoted as $L(\tau_n, B)$, that contains the root, the vertices of B and all the branch points with distances from the root to the vertices of B preserved. Figure 2 illustrates this idea with $B = \{v_1, v_2, v_3, v_4\}$; the branch points are b_1 and b_2, and in order to preserve the distances from the root to the vertices of B, the new branch from the branch points to the elements of B has a length that is equal to the sum of the edges along the path from the root to the elements of B in the original tree. Now, for a tree τ_n, we randomly

FIGURE 2. A tree on the left and its L-tree on the right corresponding to vertices $\{v_1, v_2, v_3, v_4\}$. The L-tree contains the root, the branch points b_1, b_2 and the set of vertices $\{v_1, v_2, v_3, v_4\}$. Note here that all the vertices chosen are leaves.

reorder the vertex set $\mathcal{V}(\tau_n)$ to obtain $(v_{n,1}, \ldots, v_{n,n})$. For a fixed $k < n$, we consider the L-tree of τ_n, defined by $L(\tau_n, (v_{n,1}, \ldots, v_{k,n}))$; this is akin to picking $k < n$ vertices according to a uniform distribution on $\mathcal{V}(\tau_n)$. L-trees have been used in the context of reconstruction of the trees; see Gronau and Moran [2007] for a phylogenetic applications and Aho et al. [1981] for related work in a computational context. For our purposes, we will randomly choose *only* the leaves or the terminal nodes from $\mathcal{V}(\tau_n)$ since L-trees constructed in this manner have a direct relationship with the CRT.

3.2.3. *Distributional characterisations of the CRT.* The distribution of the CRT, as the weak limit of conditioned Galton–Watson trees, remarkably can be characterised in two different ways: through strict binary subtrees that arise as a limit of the carefully chosen L-trees of conditioned Galton–Watson trees; and through a standard Brownian excursion which arises as the weak limit of suitably scaled Dyck paths of conditioned Galton–Watson trees. We first state Aldous' result for the Dyck path-based characterisation without proof.

Theorem 2. *[Aldous, 1993] Let τ_n be a conditioned Galton–Watson tree conditioned to have n vertices with an offspring distribution with mean 1 and variance $\sigma^2 \in (0, \infty)$. Let $H_n(k), 0 \leq k \leq 2n$ be the Dyck path associated with τ_n. Then, as $n \to \infty$ through the possible sizes of the unconditioned*

Galton–Watson tree,

$$\left\{\frac{1}{\sqrt{n}}H_n([2nt]), 0 \le t \le 1\right\} \Rightarrow \left\{\frac{2}{\sigma}B_t^{ex} : 0 \le t \le 1\right\}$$

where B^{ex} is the standard Brownian excursion and $\Rightarrow$ implies weak convergence of processes in $C[0,1]$, the space of continuous functions on $[0,1]$, and $[\cdot]$ stands for the integer function.

Theorem 2 implies an invariance principle with the Brownian excursion as the weak limit *regardless* of the offspring distribution as long as it is critical with finite variance; the offspring distribution manifests itself only through its standard deviation σ. This characterisation of the CRT provides a mechanism by which to study the asymptotic properties of the functionals of conditioned Galton–Watson tree models. For example, the distributions of the height of the tree, number of vertices at a given distance from the root, and total path length can be approximated by the corresponding Brownian excursion functionals. One such functional, the distance from the root of a randomly chosen vertex, will be used to develop tests in the subsequent section.

We now consider the second characterisation of the CRT using L-trees. From the construction of L-trees given earlier, it can be seen that if an L-tree of a conditioned Galton–Watson tree is constructed by choosing *only* k leaves at random, then it is a binary tree with k leaves and $2k-1$ edges. The following characterisation of the CRT is based on such binary trees $\mathcal{C}(k)$, which arise as the limit as $n \to \infty$, for a fixed k, of $L(\tau_n, \{l_1, \ldots, l_k\})$ for leaves $l_i \in \mathcal{V}(\tau_n)$.

Lemma 1. [Aldous, 1993]) There exists a consistent family $(\mathcal{C}(k), k \ge 1)$ of binary trees with k leaves which define the CRT having the density

$$f(c(k)) = \left[\prod_{i=1}^{k-1} \frac{1}{2i-1}\right]^{-1} s e^{\frac{-s^2}{2}},$$

where $s = e_1 + \cdots + e_{2k-1}$ and $c(k)$ is a binary tree with k leaves, $k-1$ branch points and $2k-1$ edges .

The factor $2^{-(k-1)}$ can be multiplied to the density if trees are considered to be ordered. The density can be used to approximate the distribution of L-trees of conditioned Galton–Watson trees from any offspring distribution with a finite variance σ^2. Aldous' proof of Lemma employed the construction from the non-homogeneous Poisson process from Section 2, and the density therefore coincides with the density in equation 1. The binary trees $\mathcal{C}(k)$ can be viewed as "marginals" of the CRT or its finite-dimensional projections. We use the density in Lemma 1 to approximate the densities of L-trees of large conditioned Galton–Watson trees τ_n, with the dependence on σ introduced in a natural manner for modelling. What is noteworthy is that through this characterisation of the CRT, not only are we able to employ the density in (1) on non-binary trees, but we also incorporate topological information into the density through the selection of the L-tree by randomly choosing a subset of the leaves.

We summarise the notations frequently used in the rest of the paper:

- τ_n: tree with n vertices including the root;
- $\tau(n)$: tree with n terminal vertices or leaves;
- $L(\tau_n, B)$: an L-tree of a tree τ_n formed by choosing a subset B of the leaves;

- σ^2: variance parameter of the offspring distribution of conditioned Galton–Watson trees.

Goodness-of-fit tests can now be developed for conditioned Galton–Watson tree models used on tree-structured data by using the CRT in two contrasting ways based on its two different characterisations.

4. Offspring variance: interpretability and estimation

As mentioned earlier, conditioned Galton–Watson trees can be equipped with a class of densities parameterized by the variance σ^2 of the offspring distribution. This parameter manifests itself as the scaling factor to the Brownian excursion from which the CRT can be constructed. It was remarked earlier that we can generate a parametric form of the density in (1) by scaling the intensity of the non-homogeneous Poisson process by a parameter θ, and by varying the value of the θ binary trees of different branch lengths but identical topology. Aldous' characterisation of the CRT through L-trees was based on the homogeneous Poisson process model, and therefore by scaling the intensity function by σ^2, we obtain a parameterised density on the limit L-trees of conditioned Galton–Watson trees,

$$f_{\sigma^2}(c(k)) = \left[\prod_{i=1}^{k-1} \frac{1}{2i-1}\right]^{-1} \frac{1}{2^{k-1}} (\sigma^2)^k s e^{\frac{-s^2\sigma^2}{2}} \qquad \sigma^2 \in \mathcal{S}, \tag{3}$$

where $c(k)$ has the edge set $\mathcal{E}(c(k)) = \{x_1, \cdots, x_{2k-1}\}$ and $\mathcal{S}$ is a suitable set of values for σ^2. The value of σ^2 in (3) has a bearing on the topology of the tree in terms of the variance of the offspring distribution or variance of the number of children at any node; different values of σ^2 lead to binary trees that differ in their topology. Since both characterisations of the CRT that are of interest in this article include σ^2, the questions of its interpretability and estimation assume importance. We first discuss how it can be estimated consistently; we then discuss some properties of parameterized densities obtained through L-trees.

4.1. Consistent estimates of σ^2.

We propose two new estimators of σ^2 based on the CRT and its relationship with the Brownian excursion. We first prove a useful result. Recall that $d(u, v)$ represented the length of the unique path connecting vertices u and v in a tree.

Proposition 4. On a conditioned Galton–Watson tree τ_n with offspring variance σ^2, suppose v is a vertex chosen according to a uniform distribution on $\mathcal{V}(\tau_n)$. Then, the random variable

$$n^{-1/2} d(\text{root}, v) \xrightarrow{d} W,$$

as $n \to \infty$, where W is a Rayleigh distributed random variable with scale $1/\sigma$.

Using Propositions 1 and 4, we can construct two different consistent estimators of σ^2, both of which are maximum likelihood estimates. The following Lemma formally states the result, the proof of which is omitted.

Lemma 2. From a random sample of conditioned Galton–Watson trees $\boldsymbol{\tau_n} = (\tau_{n_1}, \ldots, \tau_{n_p})$ from π_{σ^2}, let W_i be random variables that denote the normalised distance of a randomly chosen vertex from the root; additionally, suppose from each τ_{n_i}, choosing a subset k_i of the leaves with $|k_i|$

number of elements, an L-tree is constructed as described earlier, and its total path length S_i is obtained. Then, for fixed $|k_i|$, as, $n_i \to \infty$ and $p \to \infty$,

$$\hat{\sigma}_d^2 = \frac{2p}{\sum_{i=1}^p W_i^2} \quad \text{and} \quad \hat{\sigma}_l^2 = \frac{2\sum_{i=1}^p |k_i|}{\sum_{i=1}^p S_i}$$

are consistent estimators of $\sigma^2 > 0$.

For the case of goodness-of-fit tests based on the L-trees, σ^2 will be estimated using $\hat{\sigma}_d^2$, and for tests based on the Dyck paths, $\hat{\sigma}_l^2$ will be used for the same purpose.

4.2. **Interpretability of σ^2 in L-trees.** While the offspring variance has a clear enough interpretation as a scale for the Brownian excursion, its role in the density of L-trees is not clear. The specific question here here is the validity of $\{f_{\sigma^2}\}$ as a parametric class for L-trees of conditioned Galton–Watson trees with possible extension to the conditioned Galton–Watson trees themselves: the parameterised distribution on the L-tree needs to be extended to τ_n for every n in order to retain the interpretability of σ^2. The issue of *extendability* is crucial: suppose τ_n on n vertices contains n_l leaves and an L-tree with $k < n_l$ leaves is chosen randomly as described in Section 3.2.2. With leaves as l_i, clearly $L(\tau_n, \{l_1, \ldots, l_{n_l}\}) = \tau_n$ and it is necessary that any $f_{\sigma^2}(\cdot)$ defined on $L(\tau_n, B)$ for any $B \subset \{l_1, \ldots, l_{n_l}\}$ with $|B| = k$ can be extended to τ_n upon inserting leaves from $\{l_1, \ldots, l_{n_l}\} - B$ to $L(\tau_n, B)$, while retaining the interpretability of σ^2. Recall that $\tau_n = (\mathcal{V}(\tau_n), \mathcal{E}(\tau_n))$ resides in $\mathcal{T}_n \times \mathbb{R}_+^{n-1}$; its LCA corresponding to $B \subset \{l_1, \ldots, l_{n_l}\}$, lies in an appropriate subspace $\mathcal{T}_{2k} \times \mathbb{R}_+^{2k-1}$ with $2k < n$. Then, for all k and n with $2k < n$, the parametric density defined on $\mathcal{T}_{2k} \times \mathbb{R}_+^{2k-1}$ can be extended to $\mathcal{T}_n \times \mathbb{R}_+^{n-1}$, or is *$n$-extendable*, if $f_{\sigma^2}(\cdot)$ on $L(\tau_n, B)$ can be recovered by marginalization over $f_{\sigma^2}(\tau_n)$ on $\mathcal{T}_n \times \mathbb{R}_+^{n-1}$ for every σ^2.

Proposition 5. The class $\left\{ f_{\sigma^2} : \sigma^2 \in \mathcal{S} \right\}$ on $\mathcal{T}_{2k} \times \mathbb{R}_+^{2k-1}$is n-extendable for every n.

Remark 1. The issue of extendability was considered by Shalizi and Rinaldo [2013] in the context of exponential random graph models, where they defined a notion of the class of distributions being *projective* for the exponential family of distributions. Without going into detail, it suffices here to note that the density in (3) depends upon the branch lengths only through its sum s which, as seen in Proposition 1, is a Gamma distributed random variable. Then, $s = x_1 + \cdots + x_{2k-1}$ is the minimal sufficient statistic for σ^2 and is *separable* as defined by Shalizi and Rinaldo [2013] in the following manner: adding a leaf increases the value of the sufficient statistic by an amount equal to the length of the new edge.

5. Asymptotic tests based on CRT and its characterisations

The density in (1), the invariance principle from Theorem 2, and Lemma 1 can be used to construct goodness-of-fit tests, one- and two-sample test, for tree-structured data. This stands in contrast to its use in the binary-tree setting without the use of the CRT but purely from the non-homogeneous Poisson process. In that setting, the statistical test corresponded to a test on the Poisson process model alone. Using its relation to the CRT through L-trees, we are now in a position to develop tests to check whether the data has been generated from a conditioned Galton–Watson tree model for *any* offspring distribution with a finite variance, which as we have seen, generates

a broad variety of trees. The null hypothesis would be that the samples of trees are independent copies of a conditioned Galton–Watson tree with finite-variance offspring distribution. This greatly generalises the test in Theorem 1 which is applicable only to binary trees and represents one of the chief objectives of the article. A point requiring clarification is that the goodness-of-fit tests for trees entails testing for a generative probability model wherein the probability distribution depends on the number of vertices and branches. When we state that an independent sample is generated from a density π, we are essentially viewing them as independent trees that are generated from the same probability model, regardless of the number of vertices.

5.1. **Using L-trees.** We shall employ the estimator $\hat{\sigma}_d^2$ to construct goodness-of-fit tests for conditioned Galton–Watson trees based on L-trees. Starting with a tree τ_n with n_l number of leaves: (1) construct $\hat{\sigma_d}^2$ using the sample of ordered conditioned Galton–Watson trees; (2) for each tree, construct an L-tree by randomly choosing a subset of the leaves; (3) using the test statistic defined as the product of $\hat{\sigma_d}^2$ and the sum of the branch lengths of the L-trees, and Slutsky's theorem, construct an asymptotic rejection region based on the χ^2 distribution. The extension to two samples is straightforward. For completeness, we formally state the result.

Theorem 3. *Suppose $\boldsymbol{\tau_n} = (\tau_{n_1}, \ldots, \tau_{n_p})$ and $\boldsymbol{\eta_m} = (\eta_{m_1}, \ldots, \eta_{m_q})$ are independent samples of conditioned Galton–Watson trees from π_τ and π_η respectively with σ_τ^2 and σ_η^2 as offspring variances. Suppose that w_i and u_j are the normalised distances from the root of randomly chosen vertices from the two samples and let $\hat{\sigma_{dp}}^2 = 2p(\sum_{i=1}^p w_i^2)^{-1}$ and $\hat{\sigma_{dq}}^2 = 2q(\sum_{i=1}^q u_i^2)^{-1}$. Let $\mathbf{k} = (k_1, \ldots, k_p)$ be subsets of the leaves of τ_{n_i}, where k_i is chosen according to a uniform distribution on the leaves of τ_{n_i} and let $|k_i|$ denote the cardinality of set k_i.*

(1) For a fixed $\mathbf{k}$, define the critical function

$$\phi(\mathbf{k}, \alpha) = \begin{cases} 1 & \text{if } \frac{1}{\hat{\sigma_{dp}}^2} \sum_{i=1}^p s_i > \chi_{1-\alpha, 2\sum_{i=1}^p |k_i|}; \\ 0 & \text{otherwise,} \end{cases}$$

where s_i are the total path lengths of $L(\tau_{n_i}, k_i)$. Then, given $\mathbf{k}$, for the pair of hypotheses $H_0 : \pi_\tau = \pi$ vs $H_1 : \pi_\tau \neq \pi$, where π is the density of a conditioned Galton–Watson tree, the test given by $\phi(\mathbf{k}, \alpha)$ is such that as $n_i \to \infty$, $E_{H_0}\phi(\mathbf{k}, \alpha) \to \alpha$.

(2) Choose $\mathbf{g} = (g_1, \ldots, g_q)$ as the subset of leaves from $\boldsymbol{\eta_m}$ in a similar manner with cardinality $|g_j|$. Then, the critical function

$$\psi(\mathbf{k}, \mathbf{g}, \alpha) = \begin{cases} 1 & \text{if } \frac{\hat{\sigma_{dq}}^2 \sum_{i=1}^p s_i}{\hat{\sigma_{dp}}^2 \sum_{j=1}^q r_j} > \left(\frac{\sum_{i=1}^p |k_i|}{\sum_{j=1}^q |g_j|} \right) F_{1-\alpha, 2\sum_{i=1}^p |k_i|, 2\sum_{j=1}^q |g_j|}; \\ 0 & \text{otherwise,} \end{cases}$$

where r_j are the total path lengths of $L(\eta_{m_j}, g_j)$, defined, assuming no loss of generality that the numerator exceeds the denominator, for testing $H_0 : \pi_\tau = \pi_\eta$ against $H_1 : \pi_\tau \neq \pi_\eta$, is such that as $\min(n_i, m_j) \to \infty$ for every i, j, $E_{H_0}\psi(\mathbf{k}, \mathbf{g}, \alpha) \to \alpha$.

5.2. **Using Dyck paths.** The pertinent question when using a Dyck path to represent an ordered tree is the following: suppose a tree τ_n is distributed as a conditioned Galton–Watson tree with a certain offspring distribution; what is the ramification of the bijective transformation $\tau_n \mapsto H_n$ on the class $\{\pi_{k,\sigma^2} : k = 0, 1, 2, \ldots; 0 < \sigma^2 < \infty\}$? If we propose to develop inferential tools on

the space of Dyck paths, it is then required to establish the equivalence of statistical procedures, perhaps in the Le Cam sense, on $\{\pi_{k,\sigma^2}\}$ and the class resulting from the transformation. Indeed, this requires us to know the exact induced class prior to establishing equivalence. The probabilistic structure of the Dyck path corresponding to an *arbitrary* conditioned Galton–Watson tree is not easily ascertained; only under the special case when the offspring distribution is the geometric with success probability 1/2, is it known that the corresponding Dyck path can be modeled as a simple symmetric random walk conditioned on first return to 0 (see Aldous [1993]). This issue poses a serious difficulty if one wishes to establish some sort of equivalence between procedures on the two classes using the notion of a deficiency distance. However, weak equivalence of the procedure is easily established as a consequence of the invariance principle in Theorem 2. For a conditioned Galton–Watson tree τ_n, if $\{\mathbb{P}^n_{\sigma^2} : \sigma^2 \in \mathcal{S}\}$ is the experiment associated with its density (with respect to the counting measure) π^2_σ, then as $n \to \infty$ through the sizes of the unconditioned Galton–Watson tree, $\{\mathbb{P}^n_{\sigma^2} : \sigma^2 \in \mathcal{S}\} \Rightarrow \{\mathcal{P}^{ex}_{\sigma^2} : \sigma^2 \in \mathcal{S}\}$, where $\mathcal{P}^{ex}$ is the law on the Brownian excursion. This is our motivation in using the weak convergence argument in developing statistical models on trees: we are able to circumvent the issue of proving equivalence since the limit process is a Brownian excursion regardless of the original offspring distribution. While constructing tests, the presence of the variance parameter σ^2 as the scaling factor is handled through the use of the consistent estimator ${\hat{\sigma_l}}^2$ given in Proposition 4.

Recall that $\frac{2}{\sigma}B^{ex}$ is the limit of normalized Dyck paths that uniquely code conditioned Galton–Watson trees. Aldous' result connects the CRT to $\frac{2}{\sigma}B^{ex}$ in the following manner: Pick $U_1, \ldots, U_k$ uniformly from $[0,1]$ and consider the order statistics $U_{1:k} < \cdots < U_{k:k}$. Set $V_i = \min_{U_{i:k} \le t \le U_{i+1:k}} \frac{2}{\sigma}B^{ex}(t)$. Draw an edge of length $\frac{2}{\sigma}B^{ex}(U_{1:k})$ and label one end as the root and the other end as U_1. Inductively, from $U_{i:k}$ move back a distance $\frac{2}{\sigma}B^{ex}(U_{i:k}) - V_i$ towards the root, draw a new edge of length $\frac{2}{\sigma}B^{ex}(U_{i+1:k}) - V_i$ and label the new endpoint U_{i+1}. Aldous then proved that the resulting binary tree on k vertices with $k-1$ edges has the density given in (3). The implication of this construction is that the random tree constructed from the Brownian excursion at k uniform random times leads to the class of consistent distributions given in Lemma 1 that characterise the CRT. For a tree τ_n with n vertices, let $0 = U_{0:n} < U_{1:n} < \cdots < U_{n+1:n} = 1$ be uniform order statistics and let $V_i = \min_{U_{i:n} \le t \le U_{i+1:n}} \frac{2}{\sigma}B^{ex}(t)$. Now define the $2n+2$ dimensional vector taking values in $\mathbb{R}^{2n+2}_+$ as

$$X_n = \Big(\frac{2}{\sigma}B^{ex}(U_{i:n}), \frac{2}{\sigma}B^{ex}(V_i)\Big).$$

Based on the construction above it can be seen that the distribution of the random vector X_n defines a distribution on the random tree constructed with n vertices. One way at looking at the density in (3) via the construction above is as the density of the random variable that is the total variation of the function obtained via a linear interpolation between points in X_n; using this approach it was shown in Pitman [1999] that

$$X_n \stackrel{d}{=} \frac{\sigma(2\Gamma_{n+1})^{1/2}}{4}\Big(U_{i-1:n} - V_{i-1}, U_{i:n} - V_{i-1}; 1 \le i \le n+2 \Big| \cap_{i=1}^n (U_{i:n} > V_i)\Big),$$

where $U_{n+2:n} := 1$ and Γ_{n+1} is a Gamma random variable with shape $n+1$ and scale 1. While, in principle, it would be reasonable to define a parametric class, the distribution of X_n is not easy to compute.

To circumvent the need to compute the distribution of X_n while constructing tests, we use a random coordinate projection of B^{ex}; by this, we mean that we consider the family of distributions induced by the map $p_U : \frac{2}{\sigma}B^{ex} \mapsto \frac{2}{\sigma}B^{ex}(U)$, where U is chosen uniformly on $[0,1]$. In order for this approach to bear fruition, we first need to verify that the law on the Brownian excursion is completely determined by the law of $p_U(B^{ex})$ for a random U. In the tree setting, the approach translates to the following scenario: for a tree τ_n with offspring variance σ^2, pick a vertex v from $\mathcal{V}(\tau_n)$ according to a uniform distribution, and use the distribution of $d(\text{root}, v)$ to construct the test. This leads to the following Proposition.

Proposition 6. For an ordered conditioned Galton–Watson tree τ_n with a fixed σ^2 and Dyck path H_n, consider the class of distributions $\{r^n_{\sigma^2} : \sigma^2 \in \mathcal{S}\}$ induced through the action of choosing a vertex v at random from τ_n, and recording $n^{-1/2}H_n(s)$ for $0 \le s \le 2\sum_{i=1}^{n-1} e_i$ where s corresponds to the sum of the edges up to vertex v encountered during the depth-first walk. In similar fashion, consider the class $\{R_{\sigma^2} : \sigma^2 \in \mathcal{S}\}$ obtained from the map $p_U\left(\frac{2}{\sigma}B^{ex}\right)$ where $U \sim U[0,1]$. As $n \to \infty$, $\{r^n_{\sigma^2} : \sigma^2 \in \mathcal{S}\} \Rightarrow \{R_{\sigma^2} : \sigma^2 \in \mathcal{S}\}$, which characterises the law of $\frac{2}{\sigma}B^{ex}$.

Remark 2. The question that seeks the distribution of $d(\text{root}, v)$ given a vertex v is different from the one answered above, which is more meaningful for the following reason: the distance of a given vertex on a tree is completely determined by the value of the Dyck path that was constructed using the depth-first walk; indeed, there are several ways to uniquely code a tree and the distance of a vertex from the root should not be dictated by the choice of a traversal.

The key point is that the random variable $p_U\left(\frac{2}{\sigma}B^{ex}\right)$ is Rayleigh distributed with scale $1/\sigma$, as seen in Proposition 4. We now construct the one- and two-sample goodness-of-fit tests for conditioned Galton–Watson trees based on the Dyck path representation. Here, we use the consistent estimator $\hat{\sigma_l}^2$ from Lemma 2 to estimate σ^2. The proof is straightforward upon noting that if X is Rayleigh distributed with scale $1/b$, then X^2 is χ^2 distributed with 2 degrees of freedom and scaled by $1/b^2$.

Theorem 4. *Suppose we have two independent samples of ordered conditioned Galton–Watson trees* $\boldsymbol{\tau_n} = (\tau_{n_1}, \ldots, \tau_{n_p})$ *and* $\boldsymbol{\eta_m} = (\eta_{m_1}, \ldots, \eta_{m_q})$ *from* π_τ *and* π_η *respectively with offspring variances* σ^2_τ *and* σ^2_η. *Let* d_i *and* c_j *be the normalized distances from the root for the randomly chosen vertices as described in Proposition 4. Let* $\hat{\sigma}^2_{lp}$ *and* $\hat{\sigma}^2_{lq}$ *be the consistent estimators of offspring variances.*

(1) Define the critical function

$$\phi(\alpha) = \begin{cases} 1 & \text{if} \quad \sum_{i=1}^{p} c_i^2 > \hat{\sigma_{lp}}^2 \chi_{1-\alpha,2p}; \\ 0 & \text{otherwise}, \end{cases}$$

For the pair of hypotheses $H_0 : \pi_\tau = \pi$ *vs* $H_1 : \pi_\tau \neq \pi$, *where* π *is the density of a conditioned Galton–Watson tree, the test given by* $\phi(\alpha)$ *is such that as* $n_i, p \to \infty$, $E_{H_0}\phi(\alpha) \to \alpha$.

(2) Consider the critical function, defined assuming no loss of generality that the numerator exceeds the denominator,

$$\psi(\alpha) = \begin{cases} 1 & \text{if} \quad \frac{\hat{\sigma_{lq}}^2 \sum_{i=1}^{p} c_i^2}{\hat{\sigma_{lp}}^2 \sum_{j=1}^{q} d_j^2} > \frac{p}{q} F_{1-\alpha,2p,2q}; \\ 0 & \text{otherwise}, \end{cases}$$

For testing $H_0 : \pi_\tau = \pi_\eta$ against $H_1 : \pi_\tau \neq \pi_\eta$, the test given by $\psi(\alpha)$ is such that as $\min(n_i, m_j) \to \infty$ and $\min(p, q) \to \infty$, $E_{H_0}\psi(\alpha) \to \alpha$.

We can summarise the testing procedure as follows: (1) from each tree choose a vertex at random and compute its normalised distance from the root; (2) then, compute the estimators $\hat{\sigma}_{lp}$ and $\hat{\sigma}_{lq}$ from Lemma 2; (3) finally, determine the rejection region by computing the χ^2 or F test statistic using the normalised distances.

6. Simulations

We first briefly describe a method to simulate conditioned Galton–Watson trees by using an efficient algorithm provided by Devroye [2012] with linear expected time. This enables us to simulate a large number of conditioned Galton–Watson trees, each containing a large number of vertices—each tree is generated in expected linear time. We have made our C++ code available at `www.github.com/pkambadu/DyckPaths`. For a non-negative integer-valued random variable ξ with distribution π_k for $k = 0, 1, \ldots$, we described the construction of a Galton–Watson tree τ. However the construction of τ_n, the Galton–Watson tree conditioned to have n vertices, is not straightforward. Using a random-walk construction from n independent copies of ξ, it can be seen that in order to generate a conditioned Galton–Watson τ_n, a necessary condition is to generate a vector $\Xi = (\xi_1, \ldots, \xi_n)$ such that $\sum_{i=1}^n \xi_i = n - 1$ and then to determinine a rotation of Ξ, i.e., a vector $(\xi_k, \xi_{k+1}, \ldots, \xi_n, \xi_1, \xi_2, \ldots, \xi_{k-1})$, with the property that the total number vertices of τ equals n. Pseudo-code and a description of the algorithm, and the appropriate permutation operation that provides the final trees can be found in the Supplementary Material.

We do not examine the performance of the tests for binary trees in Theorem 1 since they are not general goodness-of-fit tests and do not capture enough topological information of the binary trees. Indeed, the test statistic is also identical to the one in Theorem 3 once the L-tree has been constructed. In the simulation exercises, our objectives are the following:

(1) simulate critical and non-critical conditioned Galton–Watson trees with various offspring distributions, and check for the accuracy of the two estimates of σ^2;
(2) examine the rejection rates for tests based on the L-trees and Dyck paths for conditioned Galton–Watson tree models, unconditioned Galton–Watson trees, and two types of models for phylogenetic trees;
(3) compare the results of the tests based on the L-trees and Dyck paths with the rejections rates obtained from Permutation tests for the corresponding test statistics.

Galton–Watson trees not conditioned on total progeny do not converge to the CRT, and we would expect the tests based on L-trees and Dyck paths to detect this. The two types of phylogenetic trees are based on a birth-death process with speciation rate of 2 on fixed taxa, and Kingman's coalescent process; these models do not correspond to conditioned Galton–Watson trees. The coalescent tree model can be used to model binary trees or dendrograms obtained from hierarchical clustering of objects. Interestingly, it has a subtle connection with the CRT [Haulk, 2011], and this manifests itself through a poor rejection rate in the tests. Since the framework proposed is based on trees with a large number of vertices, following some trials, we find that the performance of the tests is reasonable, with trees containing around 1000 vertices. Each tree can be generated in a matter

of seconds, owing to the efficiency of the algorithm provided by Devroye [2012]. Therefore, all simulation exercises are carried out with trees that contain 1000 vertices. For each tree, the branch lengths are generated from a uniform distribution between 0 and 2. Although the framework allows each tree to contain a different number of vertices, for simplicity, we ensure that each generated tree is of the same size. However, we examine the performance of the tests at different sample sizes N since we are interested in the sample sizes for consistent estimation of σ^2.

6.1. **Tests based on L-trees.** Recall from Theorem 3 that the critical function for the test was based on the statistic that corresponds to the scaled total path length or sum of branch lengths of the L-tree constructed from a randomly chosen subset of the leaves of the conditioned Galton–Watson tree; the scaling constant was estimated using the estimator $\hat{\sigma_d}^2$. Figure 3 compares the histogram of the scaled total path length of the constructed L-tree and the Gamma distribution from Proposition 1, and also reports the accuracy of the estimator; we see that the approximation is reasonable for trees with large number of vertices. In order to examine the empirical rejection rates

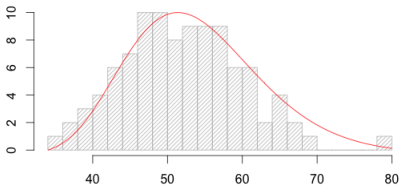

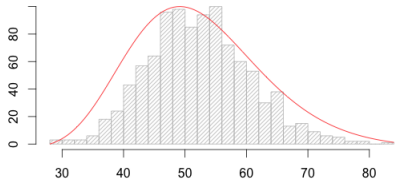

N	$\hat{\sigma_d}^2$	$\hat{\sigma_l}^2$
50	1.94	1.87
100	1.96	1.93
500	1.98	1.92
1000	2.01	1.99

FIGURE 3. Histograms of total path length, scaled by $\frac{1}{\hat{\sigma_d}^2}$, of L-trees of 100 (first histogram) and 1000 (second histogram) conditioned Galton–Watson trees from Geo(1/2) with 1000 vertices each. L-trees were constructed by randomly choosing 25 leaves. Solid red curve is the Gamma density with scale parameter 2 and shape 25. The table on the right reports the accuracy of the two estimators of the variance of Geo(1/2), $\sigma^2 = 2$, at different sample sizes.

of the test under the null and the alternative hypotheses, we generate conditioned Galton–Watson trees with varying sample sizes from some of the distributions listed in Proposition 3.1, and each containing 1000 vertices. Each simulation exercise contains 1000 simulation trials. We first compute the estimator $\hat{\sigma}_d^2$ based on choosing a vertex at random; then, we construct L-trees by randomly choosing 25 leaves in each tree; we compute the total path length of the L-trees and then compute the test statistic by scaling the total path length by $\hat{\sigma}_d^2$. We compare the performance of the χ^2 and F tests to Permutation tests, with 5000 permutations, based on the same test statistic in Table 1. The simulation results corroborate our theoretical results for goodness-of-fit tests for conditioned Galton–Watson trees. Note that the conditioned Galton–Watson tree with a Bin(0.35) distribution is based on a subcritical Galton–Watson process; however, from Proposition 3.1, we can see that in the asymptotic regime as the number of vertices grow without bound, the test in Theorem 3 still remains valid. The test is sensitive to unconditioned Galton–Watson trees as the asymptotic results are not valid for such trees; this is seen in the rejection rate for the unconditioned Galton–Watson tree with offspring distribution given by a binomial distribution with 2 trials and success probability 0.5. The trees with distribution Phylo.bd represent phylogenetic trees generated from a

	One-sample						Two-sample					
Distribution	$N = 10$		$N = 100$		$N = 1000$		$N = 10$		$N = 100$		$N = 1000$	
	χ^2	perm	χ^2	perm	χ^2	perm	F	perm	F	perm	F	perm
Geo(0.5)	0.09	0.15	0.05	0.08	0.03	0.09	0.11	0.10	0.13	0.08	0.04	0.02
Bin(0.5)	0.13	0.08	0.04	0.03	0.01	0.01	0.08	0.14	0.08	0.09	0.03	0.06
Bin(0.35)	0.10	0.16	0.12	0.07	0.06	0.08	0.21	0.14	0.13	0.10	0.04	0.01
GW-Bin(2,0.5)	0.78	0.91	0.91	0.97	0.99	1.00	0.82	0.87	0.88	0.93	0.95	1.00
Phylo.bd	0.81	0.83	0.89	0.91	0.98	0.94	0.92	0.83	0.96	0.97	0.99	1.00
Phylo.coal	0.26	0.37	0.14	0.21	0.11	0.08	0.37	0.28	0.23	0.18	0.19	0.17

TABLE 1. Tests with L-trees. One-sample: Rejection rate of one-sample χ^2 test, at $\alpha = 0.01$, in Theorem 3 and permutation test with conditioned Galton–Watson trees containing 1000 vertices under sample sizes n. GW-Bin(2,0.5) refers to an unconditioned GW tree with Bin(2,0.5) distribution; Phylo.bd and Phlyo.coal correspond to phylogenetic trees based on birth-death processes on fixed taxa and Kingman's coalescent process, respectively; Two-sample: Rejection rate of two-sample goodness-of-fit F and permutation tests, having conditioned Galton–Watson trees with offspring distribution Bin(2,0.5) under H_0 against alternatives in column 1.

birth-death process on a fixed number of taxa with speciation rate 2. These trees are not conditioned Galton–Watson trees and the test corroborates this fact. As we discussed, owing to its connection with the CRT, both tests for trees with distribution Phylo.coal, corresponding to phylogenetic trees constructed from a Kingman's coalescent process, do not seem to have good power. For binary ultrametric trees obtained from hierarchical clustering, this property will be harnessed in the data application concerning tumour MR images. We elaborate upon the connection between trees of this type and the CRT in the Discussion.

6.2. **Tests based on Dyck path.** The test statistic in this setting is the normalised distance from the root of a randomly chosen vertex which is equivalent to the value of the corresponding Dyck path at a randomly chosen point on the x-axis. Figure 4 plots the histogram of this statistic scaled by $n^{-1/2}$ where n is the number of vertices of the tree n. This offers verification of Proposition 4. An

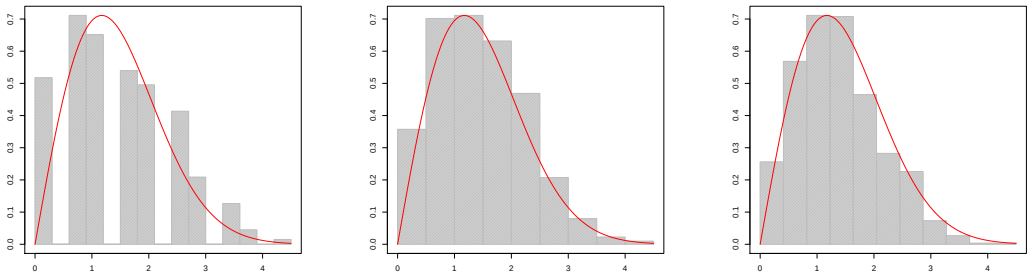

FIGURE 4. Histograms of $n^{-1/2}d(\text{root}, V)$ where V is the vertex chosen at random on 10,100 and 1000 conditioned Galton–Watson trees with 1000 vertices each from $Bin(2, 0.5)$ distribution. Solid red curve is the Rayleigh density with scale $\sqrt{2}$.

experimental setup similar to the one used to examine the performance of the tests based on L-trees provides encouraging results as seen in Table 2. The test statistic is scaled by ${\hat{\sigma_l}}^2$, the estimator of the variance obtained using an appropriate L-tree. While it may appear that the test utilises

	One-sample						Two-sample					
Distribution	$n = 10$		$n = 100$		$n = 1000$		$n = 10$		$n = 100$		$n = 1000$	
	χ^2	perm	χ^2	perm	χ^2	perm	F	perm	F	perm	F	perm
Geo(0.5)	0.14	0.22	0.07	0.07	0.04	0.01	0.18	0.09	0.11	0.08	0.05	0.03
Bin(0.5)	0.11	0.18	0.06	0.09	0.02	0.01	0.13	0.16	0.09	0.04	0.04	0.01
Bin(0.35)	0.15	0.11	0.11	0.13	0.08	0.06	0.17	0.13	0.12	0.08	0.07	0.02
GW-Bin(2,0.5)	0.68	0.83	0.88	0.91	0.94	0.99	0.84	0.81	0.93	0.95	0.93	0.96
Phylo.bd	0.89	0.78	0.94	0.96	1.00	0.99	0.85	0.93	0.94	0.92	0.99	0.99
Phylo.coal	0.46	0.66	0.23	0.22	0.18	0.13	0.28	0.31	0.21	0.19	0.15	0.09

TABLE 2. Dyck path-based tests. One-sample: Rejection rate of one-sample χ^2 test, at $\alpha = 0.01$, in Theorem 3 and permutation test with conditioned Galton–Watson trees containing 1000 vertices under sample sizes n. GW-Bin(2,0.5) refers to an unconditioned GW tree with Bin(2,0.5) distribution; Phylo.bd and Phlyo.coal correspond to phylogenetic trees based on birth-death processes on fixed taxa and Kingman's coalescent process, respectively; Two-sample: Rejection rate of two-sample goodness-of-fit F and permutation tests, having conditioned Galton–Watson trees with offspring distribution Bin(2,0.5) under H_0 against alternatives in column 1.

information of only one randomly chosen vertex of the large tree, a larger portion of the signal (tree) contributes to the test through the estimator ${\hat{\sigma_l}}^2$. Again, as in the case of L-tree-based tests, the test is powerful against unconditioned Galton–Watson trees and birth-death phylogenetic trees while performing poorly against coalescent phylogenetic trees; in fact, its power against coalescent trees is better than that of the test based on L-trees.

7. Application: Detection of tumour heterogeneity with agglomerative hierarchical clustering

We illustrate the utility of the proposed tests through a novel approach to detecting tumour heterogeneity in brain cancer by constructing binary trees obtained from MR images. We first detail the dataset, the pre-processing steps and the objective of the study; we then briefly describe the relationship between the CRT and the trees obtained from hierarchical clustering, and present results from the tests.

7.1. **Data structure, pre-processing and key scientific questions.** In this study, we used presurgical, T1-weighted post-contrast and T2-weighted/FLAIR images of 82 patients (26 women and 56 men) with histologically confirmed glioblastoma multiforme (GBM) and molecular data from The Cancer Genome Atlas (TCGA) database. (T1 and T2 indicate standard pulse sequences used in MR imaging; FLAIR is fluid attenuated inversion recovery.) The images were downloaded from The Cancer Imaging Archive and are publicly available at `https://www.cancerimagingarchive.net/`. We pre-processed the MR images and obtained 3-dimensional (3D) tumour volumes. Specifically, the images were registered spatially, followed by intensity bias correction using Medical Image Processing Analysis and Visualization software (v 6.0).13 The tumor region was segmented semi-automatically in 3D using the Medical Image Interaction Toolkit (`MITK.org`). The in-plane resolution of the image was 1mm $\times$ 1mm. In our analysis we considered T1 intensities from only the segmented regions.

The application considered here is the important problem of appropriate characterisation of tumour heterogeneity in cancer research, in the context of brain tumours. A variety of genetic, cellular

and molecular changes occur during the course of tumour development and during the course of a treatment; the characterisation of these has direct implications for the decision on suitable therapies. Methods that characterise and classify tumours based on pixel intensities are increasingly common amongst medical practitioners. The motivation behind using the vectorized pixel intensities, disregarding the spatial information, lies in the fact that if the chief objective is to characterise or classify tumour images based on heterogeneity, then the location of the heterogenous regions on the image is of little relevance: the classifying procedure needs to only *detect* the heterogeneity in the entire image in relation to another image. In this regard, the method of generating histograms of pixel intensities from magnetic resonance imaging (MRI) is particularly popular; see Just [2014] for a recent overview. The method is based on the premise that given an image from MRI, the density of pixel intensities offers a good statistical summary of the image. It has recently been shown that histogram factors such as skewness, kurtosis and percentiles are good biomarkers of tumour heterogeneity in endometrial cancer, renal cell cancer, cervical cancer and breast cancer (refer to Table 1 in Just [2014] and references within).

Our primary task is to detect group differences between patients with GBM who experience long survival times and those who experience short survival times, using the trees obtained from the tumour images. *We propose to capture the heterogeneity of the tumours by obtaining a binary tree representation of the image obtained from an agglomerative hierarchical clustering of the pixel intensities.* The resulting binary tree or dendrogram contains the individual intensities as leaves, where clusters of similarly valued intensities are the internal vertices, and branch lengths represent distances between the clusters. Such a representation expresses heterogeneity in the tumour image though the clustering of intensity values. Dendrograms, from hierarchical clustering, are obtained from images of patients that correspond to the two survival groups: patients classified as having short survival time ($<$12 months) and patients with long survival time ($\geq$ 12 months). We use agglomerative hierarchical clustering in `MATLAB`(using single, average, complete linkage algorithms) to obtain dendrograms of the image intensities. The two groups represent a 2-level classification or characterisation of tumour heterogeneity. On these two sets of binary trees we examine the performance of the two-sample F test using L-trees and Permutation tests from Theorem 3 for conditioned Galton–Watson trees with a finite offspring variance. Rejection of the null hypothesis of an equality of distribution suggests the detection of the group difference that corresponds to the heterogeneity of the tumour.

From the analysis of the rejection rates of the tests in the preceding section, we observe that the two proposed test techniques do not have good power against ultrametric trees such as coalescent trees. These trees are characterised as possessing leaves which are equidistant from the root. Such trees are commonly used as a visual representation of the results of a hierarchical clustering algorithm on objects that are based on pairwise distances. A surprising connection exists, in the asymptotic regime, between the binary trees arising from hierarchical clustering and the L-trees which characterise the CRT; see Haulk [2011] for details.

Consider the images of the patients who respectively correspond to the short (0.723 months) and long survival (57.8 months) survival times, as shown in the top panel of Figure 5. We propose to ascertain the heterogeneity between the tumour images though the topology and branch

length information of the corresponding trees. To this end, three statistics on the trees that are reasonable indicators of heterogeneity are the height of the tree, total path length, and total number of branches. The first two capture branch length information whereas the last one provides topological information. For the patient with a short survival time, these statistics, respectively, are 0.51, 3.69 and 4365; and for the patient with the much longer survival time, the corresponding statistics are 0.34, 2.13 and 3589. Intuitively, the patient with the shorter survival time, with the tumour appearing to be in an advanced stage in the image on the right in Figure 5, should have pixel intensities with richer and varied clustering tendencies; this manifests in the form of more branches but with smaller branch lengths. Our test statistic based on the scaled total path length of an L-tree constructed from the dendrogram, in this sense, appears reasonable. Observe that the leaves of the dendrogram are the individual pixel intensities, and choosing a random subset of the leaves corresponds to selecting a subset of the pixels at random without bias.

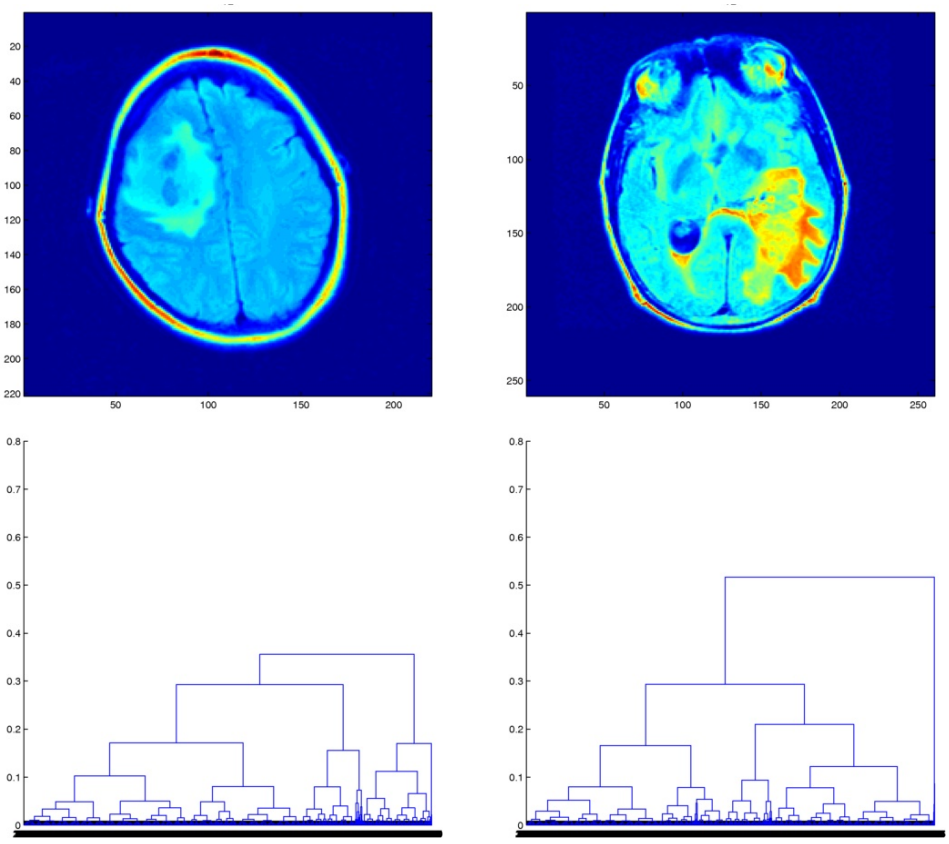

FIGURE 5. Left: Brain MR image of patient with GBM who experiences a long survival time of 57.8 months and the corresponding dendrogram (below). Right: Brain MR image of a patient with GBM who experiences a short survival time time of 0.723 months and the corresponding dendrogram (below).

7.2. **Results.** Carlsson and Mèmoli [2010] noted that the single linkage function is the most stable choice for hierarchical clustering in the sense that the distance is small between the original tree and

% of leaves chosen	Null dist.	F-statistic	Decision at 1% level
10	$F_{11638,11734}$	0.780	Reject
20	$F_{23310,23515}$	0.775	Reject
30	$F_{34793,35286}$	0.774	Reject
40	$F_{46648,47070}$	0.772	Reject

TABLE 3. At different values for the percentage of leaves chosen to construct the L-trees of the dendrograms from the images, the table reports the results of F test from Theorem 3 to test the null hypothesis of equality of distribution of the short and long survival time groups.

the tree obtained after a small perturbation of the input data; this distance is based on the Gromov–Hausdorff distance between ultrametric spaces that represent the dendrograms. We therefore used the single linkage function to measure the distance between clusters in the algorithm. Dendrograms obtained from agglomerative hierarchical clustering are sensitive to the choice of the linkage function, or the distance used between clusters. However, the large sizes of the images considered in our setting led to trees that did not appear significantly different for different choices of the linkage functions; this was corroborated by the results of the F and Permutation tests, which were unchanged in relation to different choices of linkage functions. In order to use the test prescribed in Theorem 3, we estimate the variance σ^2 with $\hat{\sigma_d}^2$. The small sample sizes—31 and 42 patients in groups who experienced short and long surviving times, respectively—do not guarantee a consistent estimate of σ^2; however, this does not appear to affect the efficacy of the tests.

Using the total path length, scaled by the estimate of the variance, of an L-tree constructed by choosing various percentages of leaves at random as the test statistic, the F test from Theorem 3 rejected the null hypothesis of equality of distribution, at 1% significance level, between the long and short survival time groups, and failed to reject H_0 when the labels were randomly permuted. Details are reported in Table 3. The Permutation test provided identical results for the same test statistic at 2000, 5000 and 10000 permutations. Importantly, the conclusions of the tests were unaltered when the linkage function was changed. Scientifically, this points to the existence of significant tumour heterogeneity among patients with different durations of survival.

8. Discussion

The circumscription of the article's attention to goodness-of-fit tests is to be viewed within the context of our efforts to initiate further research in this direction; this article is appropriately viewed as a first investigative step in the process of developing inferential tools for tree-structured data using the CRT. Aldous' papers on the CRT and variants (see, in addition, Aldous [1994a, 1991b]) provide useful distributional results and connections to common stochastic processes, which in principle can be harnessed in developing further asymptotic statistical tools. One immediate extension to this work is to approximate distributions of local tree-functionals by the corresponding Brownian excursion functionals. For example, the Wiener index of a tree, popular in phylogenetics and chemistry, is exactly $\frac{2}{n}A_n$ where A_n is the area under the curve of the Dyck path of a tree with n vertices. The distribution of A_n can be approximated by the distribution of the Brownian excursion area, which is well-know (albeit difficult to compute); see Janson [2012] and references therein for details. Other tree functionals for which the limiting distributions are approximated by known

Brownian excursion functionals include the height of a tree, which is the number of generations before extinction, and the maximal distance between a pair of vertices, among others. Preliminary investigations based on simulations appear promising. The 'weak convergence paradigm' set forth by Aldous wherein the properties (global and local) of random combinatorial objects likes trees, triangulations, and planar maps are studied through continuous approximations, offers a fertile ground for the development of statistical methodologies on such objects; see Aldous [1994b]. This is an investigative route well worth pursuing in view of the availability of large combinatorial datasets.

The absence of information on the rates of convergence of conditioned Galton–Watson trees to the CRT prevents us from obtaining a clear idea as to the size of the trees that assure the validity of our results. Our simulation experiments were conducted at varying tree sizes, and the results that validate the theory were obtained with conditioned Galton–Watson trees that contain at least 1000 vertices; the tabulated results are hence provided for trees with 1000 vertices with branch length uniformly distributed between 0 and 2. This was partly also dictated by the data application wherein the branch lengths of the dendrograms obtained from hierarchical clustering corresponded to distances between clusters, which were typically between 0 and 2. When the sample size is large, the test statistic based on L-trees can conceivably be approximated by a Normal distribution, since it a sequence of partial sums of independent random variables. But the use of the CRT provides instead leads us to our current choice.

For binary trees obtained from hierarchical clustering, which are ultrametric, we use the connection between exchangeable random hierarchies and the CRT in Haulk [2011] to develop goodness-of-fit tests. Informally stated, if the leaves of an ultrametric binary tree, $\tau(n)$ with n leaves, arising from a hierarchical clustering method are exchangeable in the sense that the distribution of the leaves is invariant to permutation or action of the symmetric group, then the L-trees $L(\tau(n), k)$ constructed by choosing k leaves converge in distribution, as $n \to \infty$, to the family $\mathcal{L}(k)$ of subtrees that characterise the CRT given in Lemma 1. A formal statement of the result, under some restrictions on the branch lengths can be found in Theorem 16, p. 72 of Haulk [2011]. As a technical note, the Theorem requires choosing the *first* k leaves based on a labelling scheme; we choose k leaves at random and suppose that they represent the first k leaves in the labelling scheme.

The data application considered in this article represents, to our knowledge, the first attempt at characterising tumour heterogeneity from images that use tree representations. However, there is a lot more to be investigated in the problem of classifying tumours based on heterogeneity: the incorporation of the spatial information of the pixels, using covariate information for each patient, and developing methods for images obtained from a longitudinal study of the patients, are all important tasks. Much remains to be done in this direction.

9. Appendix

Proof of Proposition 1. From the construction described it is easy to see that the distribution of the total path of $\tau(k)$ coincides with the distribution of the time of the kth event, say S_k. The density of S_k for a non-homogeneous Poisson process is

$$f_{S_k}(t) = e^{-m(t)} m(t)^{k-1} \lambda(t) \frac{1}{(k-1)!} \quad t \geq 0,$$

where $m(t)$ is the intensity function. Noting that when $\lambda(t) = t$, $m(t) = t^2/2$, we obtain

$$f_{S_k}(t) = e^{-t^2/2} \left(\frac{t^2}{2}\right)^{k-1} \frac{t}{(k-1)!} \quad .$$

Making a change of variable $y = t^2$, we see that $f_{S_k}(y)$ coincides with a Gamma density with shape k and scale 2. □

Proof of Proposition 3.1. Kennedy [1975] proved that under certain conditions, the study of non-critical Galton–Watson processes reduces to the study of critical ones. Specifically, if $\lambda > 0$ is a parameter such that $N_\lambda = \sum_{i\geq 0} \pi_i \lambda^i < \infty$, we set $\pi_i^\lambda = \pi_i \lambda^i / N_\lambda$. Then a conditioned Galton–Watson tree with offspring distribution π has the same distribution as a conditioned Galton–Watson tree with offspring distribution π^λ. As a consequence, if we can find a value of λ such that the conditioned Galton–Watson tree under π^λ becomes critical, then its asymptotic behaviour is similar to that in the critical case. We examine each case separately and observe that there exists a finite λ such that $N_\lambda < \infty$.

(i) Note that

$$N_\lambda = \sum_{i\geq 1} (1-p)^{i-1} p \lambda^i = \frac{p\lambda}{1-\lambda(1-p)}$$

if $(1-p)\lambda < 1$ for all p. Hence by choosing $\lambda < 1/(1-p)$ we obtain the result with the requirement that $\lambda < 1$ (only then is $\lambda < 1/(1-p)$ for all $0 < p < 1$). Here, $\lambda = 1/(2(1-p))$ results in unit mean under the new distribution.

(ii) Under this setup

$$N_\lambda = (1-p)^2 + 2p\lambda(1-p) + p^2\lambda^2$$

which is finite for all $\lambda < \infty$. Unit mean is attained at $\lambda = (1-p)/p$.

(iii) Here $N_\lambda = (1-p) + p\lambda^2$, which again is finite for all finite λ; unit mean is attained at $\lambda = \sqrt{(1-p)/p}$.

(iv) For this case $N_\lambda = p_0 + p_1\lambda + (1-p_0-p_1)\lambda^2$, which is finite for all $\lambda < \infty$ with unit mean begin attained at $\lambda = p_0/(1-p_0-p_1)$.

□

Proof of Proposition 4. Let H_n be the Dyck path that corresponds to τ_n. Then, $d(\text{root}, v)$ is distributed as $H_n(2nv)$. Since for $0 \leq s \leq 1$, $n^{-1/2} H_n(2ns)$ converges weakly in $C[0,1]$ to $B^{ex}(s)$, we can claim that $n^{-1/2} d(\text{root}, v) \xrightarrow{d} B^{ex}(v)$. Note that this conditional density of the Brownian excursion $\frac{2}{\sigma} B^{ex}$ at time $t \in (0,1)$ is given by (see Takács [1991])

$$f(t,x) = \frac{x^2\sigma^3}{4\sqrt{2\pi t^3(1-t)^3}} e^{\frac{-x^2\sigma^2}{8t(1-t)}}, \quad x > 0. \tag{4}$$

Suppose V is uniform on $[0,1]$, and we wish to determine the density of $\frac{2}{\sigma} B^{ex}(V)$, the unconditional density of $B^{ex}(V)$ as

$$r(x) = \int_0^\infty f(s,x) ds,$$

since V is uniform on $[0,1]$. Note that the map $B^{ex} \mapsto B^{ex}(V)$ is a one-dimensional random coordinate projection, and is clearly continuous on $C[0,1]$ with respect to the uniform norm. Using the continuous mapping theorem (see Billingsley [1968]), $d(\text{root}, V) \xrightarrow{d} \frac{2}{\sigma} B^{ex}(V)$, which follows a

Rayleigh distribution. □

Proof of Proposition 5

First observe that the density in (3) implies that the edge lengths are exchangeable: $s = \sum_{i=1}^{2k-1} x_i$ is invariant to permutations of x_i. This implies that relabelling the vertices and the edges leaves the distribution unchanged. For ease of notation, let

$$\mathcal{L}_k = \mathcal{T}_{2k} \times \mathbb{R}_+^{2k-1}.$$

For a tree, $\tau_n = (\mathcal{V}(\tau_n), \mathcal{E}(\tau_n))$, with n vertices and n_l number of leaves, we shall consider its L-tree, $L(\tau_n, l_1, \ldots, l_k)$, constructed from k leaves, chosen uniformly, with $k < n_l$. The question of extendability is basically a question of whether models specified in terms of joint distributions over a class of index sets are projective as defined in eq. 13, p. 92 of Kallenberg [1997]—this is the basis of our definition of n-extendability. The probability kernel p_k from $\mathcal{L}_1 \times \cdots \times \mathcal{L}_{k-1}$ to $\mathcal{L}_k$ is defined in terms of the conditional density (ignoring the normalizing factors)

$$f_{\sigma^2}\left(L(\tau_n, l_1, \ldots, l_k) | L(\tau_n, l_1, \ldots, l_{k-1})\right),$$

which is obtained as

$$\begin{aligned} f(L(\tau_n, l_1, \ldots, l_k) | L(\tau_n, l_1, \ldots, l_{k-1})) &= \frac{f(L(\tau_n, l_1, \ldots, l_k))}{f(L(\tau_n, l_1, \ldots, l_{k-1}))} \\ &\propto \frac{s'}{s} e^{-\frac{(s'^2 - s^2)\sigma^2}{2}}, \end{aligned}$$

with $s = x_1 + \cdots + x_{2k-3}$ and $s' = s + x_{2k-2} + x_{2k-1}$. By induction on k, we can extend the existence of the probability kernel p_n to $\mathcal{L}_n$ with conditional density

$$f(L(\tau_n, l_1, \ldots, l_{n_l}) | L(\tau_n, l_1, \ldots, l_{n_l - 1})) = \frac{s'}{s} e^{-\frac{(s'^2 - s^2)\sigma^2}{2}},$$

where $s = x_1 + \cdots + x_{2n_l - 3}$ and $s' = s + x_{2n_l - 2} + x_{2n_l - 1}$. By Theorem 5.17 in Kallenberg [1997], we can assert the existence of the tree τ_n with distribution $p_1 \otimes \cdots \otimes p_n$; in other words, the distribution on τ_n can be defined via the conditional densities as

$$\begin{aligned} f(\tau_n) = f(L(\tau_n, l_1)) & f(L(\tau_n, l_2) | L(\tau_n, l_1)) f(L(\tau_n, l_3) | L(\tau_n, l_1, l_2)) \\ & \ldots f(L(\tau_n, l_{n_l - 1}) | L(\tau_n, l_1, \ldots, l_{n_l})). \end{aligned}$$

Straightforward computation with the conditional densities verifies this fact. □

Proof of Proposition 6. First, we need to establish that law on the Brownian excursion is completely determined by the law of $p_u(B^{ex})$ for a random u on $[0, 1]$. For this we use a result from Cuesta-Albertos et al. [2006]. It can be checked that $m_k = \int ||x||^k ne(de)$ for $k \in \mathbb{N}$, the moments of the scaled excursion are finite, where ne is the normalized Ito measure of positive excursion of linear Brownian motion—this is true for every $\sigma^2 \in \mathcal{S}$. Now suppose ne_1 and ne_2 are two excursion measures on $\mathbb{R}_+$ and ne_1^u and ne_2^u are the randomly projected excursion measures that correspond to the projection p_u where u is uniform in $[0, 1]$. Consider the set

$$\mathbb{N}_u := \{x : ne_1^u(x) = ne_2^u(x)\}.$$

Since ne is atomless it is the case that $ne(\mathbb{E}_u) > 0$ for every u. Using the so-called Carleman condition (see, for instance, p.19 Shohat and Tamarkin [1943]) and Theorem 2.8 from Cuesta-Albertos et al. [2006], we can claim $ne_1 = ne_2$. The implication is that the distribution of the excursion is fully determined by just one random projection p_u for a uniform u in $[0, 1]$.

References

R. Adams, Z. Ghahramani, and M.I. Jordan. Tree-structured stick breaking for hierarchical data. In *NIPS*, 2010.

A . Aho, Y. Sagiv, T. G. Szymanski, and J. D. Ullman. Inferring a Tree from Lowest Common Ancestors with an Application to the Optimization of Relational Expressions. *SIAM Journal of Computation*, 10:405–421, 1981.

D. Aldous. The Continuum Random Tree I. *Annals of Probability*, 19:1–28, 1991a.

D. Aldous. Asymptotic Fringe Distributions for General Families of Random Trees. *Annals of Applied Probability*, 1:228–266, 1991b.

D. Aldous. The Continuum Random Tree III. *Annals of Probability*, 21:248–289, 1993.

D. Aldous. Recursive Self-similarity for Random Trees, Random Triangulations and Brownian Excursion. *Annals of Probability*, 22:527–545, 1994a.

D. Aldous. Triangulating the Circle, at Random. *The American Mathematical Monthly*, 101:223–233, 1994b.

D. Aldous. Probability Distributions on Cladograms. *In: Random Discrete Structures, ed. David Aldous and R. Pemantle, IMA Volumes Math. Appl*, 76:1–18, 1996.

B. Aydin, G. Pataki, H. Wang, E. Bullitt, and J. S. Marron. A Principal Component Analysis for Trees. *Annals of Applied Statistics*, 3:1597–1615, 2009.

B. Aydin, G. Pataki, H. Wang, A. Ladha, E. Bullitt, and J. S. Marron. Visualizing the Structure of Large Trees. *Electronic Journal of Statistics*, 5:405–420, 2011.

S. Aylward and E. Bullitt. Initialization, Noise, Singularities and Scale in Height Ridge Traversal for Tubular Object Centerline Extraction. *IEEE Transactions of Medical Imaging*, 21:61–75, 2002.

P. Billingsley. *Weak Convergence of Probability Measures*. John Wiley and Sons, New York, 1968.

J. R. Busch, P. A. Ferrari, A .G. Flesia, R. Fraiman, S. P. Grynberg, and F. Leonardi. Testing statistical hypothesis on random trees and applications to the protein classification problem. *Annals of Applied Statistics*, 3:542–563, 2009.

G. Carlsson and F. Mèmoli. Characterization, Stability and Convergence of Hierarchical Clustering Methods. *Journal of Machine Learning Research*, 11:1425–1470, 2010.

J. A. Cuesta-Albertos, R. Fraiman, and T. Ransford. Random projections and goodness-of-fit tests in infinite-dimensional spaces. *Bulletin of Brazilian Mathematical Society, New Series*, 37:1–25, 2006.

L. Devroye. Simulating Size-constrained Galton-Watson Trees. *SIAM Journal of Computing*, 41: 1–11, 2012.

I. Gronau and S. Moran. Neighbor Joining Algorithms for Inferring Phylogenies via LCA Distances. *Journal of Computational Biology*, 14:1–15, 2007.

P. Guttorp. *Statistical Inference for Branching Processes*. Wiley Series in Probability and Statistics, 1991.

C. Haulk. *Aspects of Exchangeable Partitions and Trees*. Ph.D. Dissertation, University of California, Berkeley., 2011.

K. A. Heller and Z. Ghahramani. Bayesian hierarchical clustering. In *ICML*, 2005.

Y. Hu, J. L. Boyd-Garber, H. Daume, and Z. I. Ying. Binary to bushy: Bayesian hierarchical clustering with the beta coalescent. In *NIPS*, 2013.

S. Janson. Simply generated trees, conditioned GaltonWatson trees, random allocations and condensation. *Probability Surveys*, 9:103–252, 2012.

N. Just. Improving tumour heterogeneity MRI assessment with histograms. *British Journal of Cancer. doi : 10.1038/bjc.2014.512*, pages 1–9, 2014.

O. Kallenberg. *Foundations of Modern Probability Theory, First edition* . Springer-Verlag, New York, 1997.

D. P. Kennedy. The Galton-Watson Process Conditioned on the Total Progeny. *Journal of Applied Probability*, 12:800–806, 1975.

D. L. Kolbe and S. R. Eddy. Local RNA structure alignment with incomplete sequence. *Bioinformatics*, 25:1236–1243, 2009.

S. Y. Le, R. Nussinov, and J. V. Maizel. Tree graphs of RNA secondary structures and their comparisons. *Computers and Biomedical Research*, 22:461–473, 1989.

R. Neal. Defining Priors for Distributions Using Dirichlet Diffusion Trees. *Tech. Report No. 0104, Dept. Statistics, University of Toronto*, pages 1–25, 2001.

J. Pitman. Brownian Motion, Bridge, Excursion, and Meander Characterized by Sampling at Independent Uniform Times. *Electronic Journal of Probability*, 4:1–33, 1999.

J. Pitman. *Combinatorial Stochastic Processes*. Springer-Verlag, New York, 2006.

P. S. La Rosa, B. Shands, E. Deych, Y. Zhou, E. Sodergen, G. Weinstock, and W D Shannon. Statistical Object Data Analysis on Taxonomic Trees from Human Microbiome Data. *PLoS ONE*, 7:1–12, 2012.

W. R. Schmitt and M. S. Waterman. Linear trees and RNA secondary structure. *Discrete Applied Mathematics*, 51:317–323, 1994.

C. R. Shalizi and A. Rinaldo. Consistency under Sampling of Exponential Random Graph Models. *Annals of Statistics*, 41:508–535, 2013.

D. Shen, H. Shen, S. Bhamidi, Y. Munoz-Maldonado, Y. Kim, and J. S. Marron. Functional Data Analysis of Tree Structured Objects. *Journal of Computational and Graphical Statistics*, DOI:10.1080/10618600.2013.786943, 2013.

J. A. Shohat and J. D. Tamarkin. *The Problem on Moments*. 1943.

J. M. Steele. Gibbs' Measures on Combinatorial Objects and The Central Limit Theorem For an Exponential Family of Random Trees. *Probability in the Engineering and Informational Sciences*, 1:47–59, 1987.

L. Takács. A Bernoulli Excursion and Its Various Applications. *Advances in Applied Probability*, 23:557–585, 1991.

S. Tatikonda and S. Parthasarathy. Hashing tree-structured data: Methods and applications. In *IEEE 26th International Conference on Data Engineering, Long Beach*, 2010.

Y. W. Teh, H. Daume, and D. M. Roy. Bayesian agglomerative clustering with coalescents. In *NIPS*, 2008.

H. Wang and J. S. Marron. Object oriented data analysis: Sets of trees. *Annals of Statistics*, 35: 1849–2311, 2007.

Y. Wang, J. S. Marron, B. Aydin, A. Ladha, E. Bullitt, and H. Wang. A Nonparametric Regression Model with Tree-structured Response. *Journal of the American Statistical Association*, 107: 1272–1285, 2012.

R. Yang, P. Kalnis, and A. K. H. Tung. Similarity evaluation on tree-structured data. In *SIGMOD, Baltimore*, 2005.

School of Mathematical Sciences, University of Nottingham
University Park, Nottingham, NG7 2RD, U.K.
E-mail address: karthik.bharath@nottingham.ac.uk

Bloomberg Labs
New York, USA
E-mail address: pkambadur@bloomberg.net

Department of Statistics, University of Connecticut
215 Glenbrook Road, Storrs, CT 06269, USA
E-mail address: dipak.dey@uconn.edu

Department of Bioinformatics and Computational Biology
The University of Texas MD Anderson Cancer Center
1515 Holcombe Blvd, Houston, TX 77030, USA
E-mail address: aruppore@mdanderson.org

Department of Biostatistics, The University of Texas MD Anderson Cancer Center
1515 Holcombe Blvd, Houston, TX 77030, USA
E-mail address: veera@mdanderson.org